1

What Is Artificial Intelligence?

Artificial intelligence (AI) consists of machines that are as smart, or smarter, than humans. Most often, computers control these machines. There are many examples of artificial intelligence in science fiction films and shows. One example is the Cylons in the TV and film franchise *Battlestar Galactica,* which are robots that turned against humans. Other examples include the smart robots and Skynet from the *Terminator* movies and TV series. These machines were examples of evil artificial intelligence that were trying to kill people. If this is what AI research represents, do we really want to make intelligent machines?

On the other hand, everybody who watches *Star Trek* knows that artificial intelligence can be a good thing. The **android** character Lieutenant Commander Data uses his intelligence and strength to help his crewmates. Because of his design, he will not hurt people. If intelligent robots of the future are made like Data, programmed not to harm humans, we should not worry. Instead, we should be developing AI as quickly as possible.

The previous examples make opinions about AI seem simple: If smart robots are evil, then we should not create them; if smart robots are good, then perhaps we should create them. However, nothing is ever that simple. For example, as sci-fi fans know, Data had an evil "brother," and there were a few Cylons and Terminator

Figure 1.1 Lieutenant Commander Data (actor Brent Spiner) checks out B-4, his less-advanced predecessor, in the 2002 film *Star Trek: Nemesis*.

robots that were good. They were actually helpful to humans. Even in these popular fictional shows and movies, the truth is that intelligent machines were neither completely good nor bad. Today we are unclear as to whether artificial intelligence is going to end up as humanity's savior, destroyer, or simply another tool for us to use. In any case, you need to know the basics about AI before deciding on its pros and cons, so keep reading.

WHAT ARE SENTIENCE AND INTELLIGENCE?

Two key words used a lot in AI research are ***sentience*** and ***intelligence***. The first term, *sentience*, is a little easier to understand.

Sentience is the ability to feel and to experience perceptions—in other words, to experience the world through the senses of sight, hearing, touch, smell, and taste. How can computers or robots be sentient? A computer can be hooked up to a camera or a scanner that feeds images into it the same way that human eyes feed images into the brain. Thus, in this sense, computers can see. Also, we can hook up computers to microphones and even program them to recognize words when a person talks into a microphone or telephone. Therefore, computers can also hear. Incredibly, there are even devices that help computers to smell, taste, and touch, although these are still being developed and perfected. The bottom line is that some computers and robots are able to sense the world. Whether or not they can understand or appreciate what their senses tell them is something else.

Some scientists think that there is more to sentience than just being able to sense the world. They think that to be truly sentient one must be able to experience the sensation in an emotional way. When a dog smells food, it gets excited. When it sees a family member, it becomes happy. These scientists would say that even after attaching a camera to a computer to allow it to see, the computer would not feel excited if its owner walked in front of the lens. There is a difference between a computer "seeing" a scanned image and a dog or a person seeing something that makes them happy. By definition, we might say that a computer is sentient because it can be fitted with cameras, microphones, and other sensors to touch and taste. However, some scientists would still say that just because a computer can sense, it is not necessarily true that it is sentient, or sensitive in perception or feeling.

The other key word in AI research is *intelligence*. It is difficult to precisely define this word. An intelligent person is smart, and intelligence has something to do with the ability to think. A dictionary definition of *intelligence* might say something like, "having the ability to learn new things, to learn from experience, and to use this knowledge to solve problems." Let's try to put this definition to use in AI.

Part of the definition of *intelligence* is about the ability to learn, and most of us would agree that learning means that we are adding new knowledge about something to our brains. When students are young, they learn that 2 + 2 = 4, and they learn that George

Washington was the first U.S. president. Before they learned these facts, they did not know them. Learning these things added knowledge to their brains. So how does that work with a machine? If a person loads his or her computer with the *Encyclopedia Britannica,* has the computer learned everything in the encyclopedia? If so, what happens when that information is deleted? Has the computer now "unlearned," or forgotten, the encyclopedia? Or, is there more to learning than simply adding facts to a person's brain?

Another part of the definition of intelligence is the ability to learn from experience. For example, say a person turns off a light bulb and grabs it right away. After burning his fingers, he learns that this is not a good idea. He might even learn to wait a few minutes to let the light bulb cool off, or to use a towel or washcloth to keep his fingers from burning. The point is that this person learns from the bad experience of grabbing a hot light bulb.

Computers can learn, too. People can take a computer, robot, or any other device that has computer controls and fit it with temperature sensors and a mechanical arm. Then, we can let it grab a number of different objects, programming it to recognize when something is so hot it would hurt a person. We could program it not to pick up the object again in the future. Thus, the experience of picking up a hot light bulb—even if the computer does not feel pain—can still teach it to avoid hot light bulbs in the future. In other words, we can say that the computer has learned from its experience of picking up a hot light bulb. This means that computers can pass two of the tests for intelligence mentioned earlier. However, what about the third test? Can a computer use what it has learned to solve problems?

To some extent, computers can learn to solve problems, depending on the sorts of problems they receive. For example, a computer can use its thermal sensors and programming to solve the simple problem of how long it must wait after turning off a light bulb to safely pick it up. On the other hand, a computer probably cannot fight a war all on its own. That sort of problem is too complicated, since there are too many jobs to do and too many decisions for a computer to make. Computers are even worse at recognizing right from wrong, especially when the issues are not clear-cut. If two people are arguing about who should get a car, a computer might simply decide that each person should get half. Although computers might be able to use their experiences to solve simple problems, they cannot handle the more complicated ones people deal with every day. This does not mean that computers will *never* be able to solve these sorts of hard problems. It just means they can't do it yet.

MAKING MACHINES THAT CAN MAKE "INTELLIGENT" DECISIONS

Part of having intelligence is being able to make decisions. Some easy decisions might include whether or not to go to the restroom, or if it is safe enough to cross the street. People make thousands of decisions every day. In fact, any animal with a brain can make a decision of sorts—depending on how we define the word *decision*. An ant crawling across the ground might run into a stick and be forced to "decide" whether to crawl over it, turn right, or turn left, but an ant's "decisions" are made by following very simple rules that are encoded into its DNA. Thus, an ant might "decide" to turn right when it runs into a stick, but this sort of decision is sort of similar to saying that a ball "decides" to drop when it rolls off the end of a table. The ant has no more choice over which way it turns than a ball has a choice about falling—something instinctive like this is probably not what we think of as being a decision. A dog might be able to choose between a tasty steak and dry dog food, but it can't decide what movie it would like to see on TV. The decisions that people make are more complicated than the ones made by ants or dogs because human lives are more complicated than the lives of animals. The reason that humans can have complicated lives is that our intelligence has enabled us to create a more complicated world.

Before most people make decisions, they consider the pros and cons of their options. Their choices come from *all* of the information available. There are, however, other ways of making decisions that don't require as much thought. One example of this is a decision made several times each day on nuclear submarines. Submarines cannot hear very well what's behind them, so every so often they engage in maneuvers to ensure no one is following them.

There are several ways of finding out what might be behind them, including the use of baffle-clearing patterns. A submarine captain doesn't want to be predictable. If the submarine's actions are expected, an enemy will be able to hunt the vessel. Therefore, the control room of every submarine contains a pair of dice. Every so often, one of the officers will roll the dice two times. The first roll decides when the submarine will conduct its maneuver. The second roll decides the pattern chosen by the captain, who consults a table such as the one shown here. By allowing the dice to make these decisions, the captain knows that his decisions will be purely random. An enemy captain can't predict what will happen or when it will happen.

This is the kind of decision a computer can easily make. The only problem is that it is not an *intelligent* decision. Chance makes the choice, rather than a thinking process. Sometimes random decisions are a good way to go, but they can't be confused with any decision that requires intelligence.

Table 1.1: Submarine "Baffle-clearing" Patterns

Dice Roll	Time of maneuver	Maneuver to use
1	Between 00 and 05 (on the hour to 5 minutes after the hour)	Pattern 1
2	05-10	
3	10-15	Pattern 2
4	15-20	
5	20-25	Pattern 3
6	25-30	
7	30-35	Pattern 2
8	35-40	
9	40-45	Pattern 1
10	45-50	
11	50-55	Pattern 3
12	55-00	

Table 1.1 This is an example of a table that a submarine captain might use to decide which pattern to adopt when checking to see if he is being followed.

What Is IQ?

For thousands of years, people have recognized that some people are smarter than others. Yet they had trouble measuring this. What was needed was a way to measure intelligence. In 1912, a German psychologist named William Stern was the first to use the term *intelligence quotient*, abbreviated *IQ*. This quotient was thought to be a way to measure a person's intelligence. Two other scientists, Alfred Binet and Theodore Simon, developed the first IQ test. To this day, however, not everyone agrees on what this exam and its results really mean.

At first glance, a person's IQ score is supposed to measure his or her intelligence. A person of average intelligence will have an IQ of 100. A score lower than 100 means an intelligence that is below average. Scores higher than 100 signify an intelligence that is higher than average. This is not a bad assumption. The problem is that measuring human intelligence is a lot more complex than people originally imagined.

Scientists have come to realize that there are many different kinds of intelligence. Some people are good at math, while others excel at writing. There are people who are amazing in music, art, or recognizing patterns. Judging people's intelligence based on how well they do on a math and English test might not be fair if they have stronger abilities in other skills. Thus, giving everyone the same test might not be the best way to determine who is smarter, or who has a higher IQ.

This leads us to an interesting question: If we can't understand how to measure intelligence in people, then how are we going to know how intelligent machines are?

Computers are also good at making simple decisions, similar to the way an ant does. An ant's decisions are rooted in answers embedded in its DNA. Computers can become programmed with rules telling them what to do in common situations. People, on the other hand, do not use simple rulebooks while making decisions. There are many decisions that need more than math or simple rules, such as buying a car. Some people like to buy cars that will impress other people. However, a computer doesn't know that a Mercedes might impress people in one neighborhood, while a Volvo might rank higher in another area. Other people simply fall in love with a certain kind of car and want to buy it for that reason. Computers don't know how to tell the difference between a car that a teen might love and one that a parent might love.

Every day, people decide whether to get married, buy a home, or change to a profession that will make them happier. Some of the most important decisions people make, however, are the most challenging for a computer. Of course, some of these decisions are not easy for people to make either. When people let machines make complicated decisions, the machines often make mistakes. For

How Do We Measure a Machine's Intelligence?

Measuring the intelligence of a machine is more difficult than figuring out the intelligence of a person. We can program a machine to know the answers to almost any factual question. Computers are better than humans at figuring out math problems and memorizing grammar rules. If a teacher gives a computer a math or science test, the computer will look like the smartest "person" taking the exam. However, if the teacher asks a computer to write an essay, it will falter. If the teacher asks it to find symbolism in a painting, it will do worse than most people will.

Many of the intelligence or problem-solving tests that researchers give to people are not good ways to test computers. For example, math problems are used to test human intelligence, but computers can be programmed to do perfectly on math questions. Tests that rely on facts are not helpful in determining if a computer is intelligent, because facts can be loaded into a computer's hard drive. In those sorts of difficult decisions, there is still no substitute for a human mind.

MAKING MACHINES THAT CAN THINK

Computers can become sentient, they can learn, and they can make simple decisions. Yet, in spite of this, computers today are not yet what one would call intelligent. All of this raises a question: How do scientists make a machine that can actually *think*? How do they make a computer that can solve complicated problems that call for more than simple, random chance or simple rules?

In humans, IQ measures intelligence. Human intelligence resides in the brain, which is made of cells called *neurons*. Some computer scientists think that brain cells serve the same function as the

1950, one of the first computer scientists, Alan Turing, determined the best test of a machine's intelligence. The test should not consist of asking a machine to answer specific questions. Rather, a person should try to carry on a conversation with the machine. Turing suggested putting a human judge in a room with two terminals: one with a person on the other end and the other hooked up to a computer. The judge wouldn't know which terminal had the person on the other end and which had the computer. The judge's job was to try to figure out which one was the computer and which was the person. If a computer can fool the judge into thinking it's a person, then the computer is intelligent.

One problem with the Turing test is that we don't know if a computer is really thinking or if it is just very cleverly simulating thought. We can only know what the computer communicates with us. On the other hand, we can say the same thing about everyone around us. When we talk with another person, how can we know what is actually taking place inside his or her head? The Turing test might have some potential problems, but it is still about the best we can do.

transistors on a computer chip. These researchers feel that human intelligence depends both on the number of neurons and on how these cells are connected. The average human brain contains 100 billion neurons. This is 333 times more neurons than there are in a cat's brain, and 1 million to 10 million times more than in an ant's brain.

According to these scientists, the reason humans are intelligent is due to all the neurons in the human brain. They believe computers could become as intelligent as humans are if the function and number of computer chip transistors were the same as brain neurons. One problem with this approach is that examples exist that do not support it. For instance, whales and elephants have twice as many neurons as humans and are very smart. However, research shows they are not as intelligent as humans are. If that's the case, then intelligence must depend on more than just the number of brain cells or the number of transistors. Today's high-performance computers have about a billion transistors per chip. This number is close to the amount of neurons a dog brain has. Yet most people would agree that computers and dogs do not have the same level of intelligence. This suggests that there is more to intelligence than just having the right **hardware**. If scientists are going to try to develop artificial intelligence, then they must figure out how smart a machine is.

Other scientists have pointed out that there is more to our brains than just the number of neurons. Each neuron sends out tens or hundreds of stalks, or *dendrites*, which touch other neurons. With time, these linkages change. For instance, when a person learns something new, his or her neurons rewire themselves to connect differently than they had before learning the new information. The brain is a dynamic network of interconnected cells. Maybe, then, the *complexity* of the brain (the neurons *plus* their connectedness) is what makes people intelligent. If that's the case, scientists need to make a computer that has a huge number of transistors, each of which is connected with tens or hundreds of other transistors on the same computer chip. This has not yet been accomplished.

There are two main schools of thought among AI researchers. One camp believes that intelligence will automatically appear when the hardware becomes complicated enough. A computer will become intelligent when it has as many transistors as a human brain has cells. These transistors must be interconnected, like brain cells. Other scientists argue that this is a little too simple. They think that it will take not only hardware, but also **software**. If this is the case, then the hardware will need to be complex enough to run the program. Yet the programming is what makes the computer intelligent.

The scientists who feel that software is the key to artificial intelligence are trying to write intelligent programs that will help the computer to learn, use what it has learned, and make decisions. In AI research, this is the basic issue that needs to be overcome. Unless scientists can answer this question, the only way they're going to design an intelligent machine is through sheer luck. Today's common belief is that the making of an intelligent machine requires both first-rate hardware and cutting-edge software.